Q-PROP: SAMPLE-EFFICIENT POLICY GRADIENT WITH AN OFF-POLICY CRITIC

Shixiang Gu[123], **Timothy Lillicrap**[4], **Zoubin Ghahramani**[1], **Richard E. Turner**[1], **Sergey Levine**[35]
sg717@cam.ac.uk, countzero@google.com, zoubin@eng.cam.ac.uk,
ret26@cam.ac.uk, svlevine@eecs.berkeley.edu
[1]University of Cambridge, UK
[2]Max Planck Institute for Intelligent Systems, Tübingen, Germany
[3]Google Brain, USA
[4]DeepMind, UK
[5]UC Berkeley, USA

arXiv:1611.02247v2 [cs.LG] 17 Nov 2016

ABSTRACT

Model-free deep reinforcement learning (RL) methods have been successful in a wide variety of simulated domains. However, a major obstacle facing deep RL in the real world is the high sample complexity of such methods. Unbiased batch policy-gradient methods offer stable learning, but at the cost of high variance, which often requires large batches, while TD-style methods, such as off-policy actor-critic and Q-learning, are more sample-efficient but biased, and often require costly hyperparameter sweeps to stabilize. In this work, we aim to develop methods that combine the stability of unbiased policy gradients with the efficiency of off-policy RL. We present Q-Prop, a policy gradient method that uses a Taylor expansion of the off-policy critic as a control variate. Q-Prop is both sample efficient and stable, and effectively combines the benefits of on-policy and off-policy methods. We analyze the connection between Q-Prop and existing model-free algorithms, and use control variate theory to derive two variants of Q-Prop with conservative and aggressive adaptation. We show that conservative Q-Prop provides substantial gains in sample efficiency over trust region policy optimization (TRPO) with generalized advantage estimation (GAE), and improves stability over deep deterministic policy gradient (DDPG), the state-of-the-art on-policy and off-policy methods, on OpenAI Gym's MuJoCo continuous control environments.

1 INTRODUCTION

Model-free reinforcement learning is a promising approach for solving arbitrary goal-directed sequential decision-making problems with only high-level reward signals and no supervision. It has recently been extended to utilize large neural network policies and value functions, and has been shown to be successful in solving a range of difficult problems (Mnih et al., 2015; Schulman et al., 2015; Lillicrap et al., 2016; Silver et al., 2016; Gu et al., 2016b; Mnih et al., 2016). Deep neural network parametrization minimizes the need for manual feature and policy engineering, and allows learning end-to-end policies mapping from high-dimensional inputs, such as images, directly to actions. However, such expressive parametrization also introduces a number of practical problems. Deep reinforcement learning algorithms tend to be sensitive to hyperparameter settings, often requiring extensive hyperparameter sweeps to find good values. Poor hyperparameter settings tend to produce unstable or non-convergent learning. Other of these algorithms tend to exhibit high sample complexity, often to the point of being impractical to run on real physical systems. Although a number of recent techniques have sought to alleviate some of these issues (Hasselt, 2010; Mnih et al., 2015; Schulman et al., 2015; 2016), these recent advances still provide only a partial solution to the instability and sample complexity challenges.

Model-free reinforcement learning consists of on- and off-policy methods. Policy gradient methods (Peters & Schaal, 2006; Schulman et al., 2015) are popular on-policy methods that directly maximize the cumulative future returns with respect to the policy. While these algorithms use unbi-

ased gradient estimators of the true reinforcement learning objective, their estimators rely on Monte Carlo returns and have high variance. To cope with high variance gradient estimates and difficult optimization landscapes, a number of techniques have been proposed, including constraining the change in the policy at each gradient step (Kakade, 2001; Peters et al., 2010) and mixing value-based back-ups to trade off bias and variance in Monte Carlo return estimates (Schulman et al., 2015). However, these methods all tend to require very large numbers of samples to deal with the high variance when estimating gradients of high-dimensional neural network policies. The crux of the problem with policy gradient methods is that they can only effectively use on-policy samples, which means that they require collecting large amounts of on-policy experiences after each parameter update to the policy. This makes them very sample intensive. Off-policy methods, such as Q-learning (Watkins & Dayan, 1992; Sutton et al., 1999; Mnih et al., 2015; Gu et al., 2016b) and off-policy actor-critic methods (Lever, 2014; Lillicrap et al., 2016), can instead use all samples, including off-policy samples, by adopting temporal difference learning with experience replay. Such methods are much more sample-efficient. However, convergence of these algorithms is in general not guaranteed with non-linear function approximators, and practical convergence and instability issues typically mean that extensive hyperparameter tuning is required to attain good results.

In order to make deep reinforcement learning practical as a tool for tackling real-world tasks, we must develop methods that are both data efficient and stable. In this paper, we propose Q-Prop, a step in this direction that combines the advantages of on-policy policy gradient methods with the efficiency of off-policy learning. Unlike prior approaches for off-policy learning, which either introduce bias (Sutton et al., 1999; Silver et al., 2014) or increase variance (Precup, 2000; Levine & Koltun, 2013; Munos et al., 2016), Q-Prop can reduce the variance of gradient estimator without adding bias; unlike prior approaches for critic-based variance reduction (Schulman et al., 2016) which fit the value function on-policy, Q-Prop learns the action-value function off-policy. The core idea is to use the first-order Taylor expansion of the critic as a control variate, resulting in an analytical gradient term through the critic and a policy gradient term consisting of the residuals in advantage approximations. The method helps unify policy gradient and actor-critic methods: it can be seen as using the off-policy critic to reduce variance in policy gradient or using on-policy Monte Carlo returns to correct for bias in the critic gradient. We further provide theoretical analysis of the control variate, and derive two additional variants of Q-Prop. The method can be easily incorporated into any policy gradient algorithm. We show that Q-Prop provides substantial gains in sample efficiency over trust region policy optimization (TRPO) with generalized advantage estimation (GAE) (Schulman et al., 2015; 2016), and improved stability over deep deterministic policy gradient (DDPG) (Lillicrap et al., 2016) across a repertoire of continuous control tasks.

2 BACKGROUND

Reinforcement learning (RL) aims to learn a policy for an agent such that it behaves optimally according to a reward function. At a time step t and state $\boldsymbol{s}_t$, the agent chooses an action $\boldsymbol{a}_t$ according to its policy $\pi(\boldsymbol{a}_t|\boldsymbol{s}_t)$, the state of the agent and the environment changes to new state $\boldsymbol{s}_{t+1}$ according to dynamics $p(\boldsymbol{s}_{t+1}|\boldsymbol{s}_t,\boldsymbol{a}_t)$, the agent receives a reward $r(\boldsymbol{s}_t,\boldsymbol{a}_t)$, and the process continues. Let R_t denote a γ-discounted cumulative return from t for an infinite horizon problem, i.e $R_t = \sum_{t'=t}^{\infty} \gamma^{t'-t} r(\boldsymbol{s}_{t'},\boldsymbol{a}_{t'})$. The goal of reinforcement learning is to maximize the expected return $J(\theta) = \mathbb{E}_{\pi_\theta}[R_0]$ with respect to the policy parameters θ. In this section, we review several standard techniques for performing this optimization, and in the next section, we will discuss our proposed Q-Prop algorithm that combines the strengths of these approaches to achieve efficient, unbiased RL.

2.1 POLICY GRADIENT METHODS

Policy gradient methods[1] apply direct gradient-based optimization to the reinforcement learning objective. This involves directly differentiating the $J(\theta)$ objective with respect to the policy parameters θ. The standard form, known as the REINFORCE algorithm (Williams, 1992), is shown below:

$$\nabla_\theta J(\theta) = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)\gamma^t R_t] = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t \nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)(R_t - b(\boldsymbol{s}_t))], \tag{1}$$

[1] "Policy gradient" technically refers to a more general class of methods (Sutton et al., 1999); however, to make it compatible with references in other work we use it to primarily refer REINFORCE (Williams, 1992).

where $b(\boldsymbol{s}_t)$ is known as the baseline. For convenience of later derivations, Eq. 1 can also be written as below, where $\rho_\pi(\boldsymbol{s}) = (1-\gamma)\sum_{t=0}^{\infty}\gamma^t p(\boldsymbol{s}_t = \boldsymbol{s})$ is the normalized state visitation frequency,

$$\nabla_\theta J(\theta) = \mathbb{E}_{\boldsymbol{s}_t\sim\rho_\pi(\cdot),\boldsymbol{a}_t\sim\pi(\cdot|\boldsymbol{s}_t)}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)(R_t - b(\boldsymbol{s}_t))]. \tag{2}$$

The gradient is estimated using Monte Carlo samples in practice and has very high variance. A proper choice of baseline is necessary to reduce the variance sufficiently such that learning becomes feasible. A common choice is to estimate the value function of the state $V_\pi(\boldsymbol{s}_t)$ to use as the baseline, which provides an estimate of advantage function $A_\pi(\boldsymbol{s}_t)$, a centered action-value function $Q_\pi(\boldsymbol{s}_t)$, each defined below,

$$\begin{aligned} V_\pi(\boldsymbol{s}_t) &= \mathbb{E}_\pi[R_t] = \mathbb{E}_{\pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)}[Q_\pi(\boldsymbol{s}_t,\boldsymbol{a}_t)] \\ Q_\pi(\boldsymbol{s}_t,\boldsymbol{a}_t) &= r(\boldsymbol{s}_t,\boldsymbol{a}_t) + \gamma\mathbb{E}_\pi[R_{t+1}] = r(\boldsymbol{s}_t,\boldsymbol{a}_t) + \gamma\mathbb{E}_{p(\boldsymbol{s}_{t+1}|\boldsymbol{s}_t,\boldsymbol{a}_t)}[V_\pi(\boldsymbol{s}_{t+1})] \\ A_\pi(\boldsymbol{s}_t,\boldsymbol{a}_t) &= Q_\pi(\boldsymbol{s}_t,\boldsymbol{a}_t) - V_\pi(\boldsymbol{s}_t). \end{aligned} \tag{3}$$

$Q_\pi(\boldsymbol{s}_t)$ summarizes the performance of each action from a given state, assuming it follows π thereafter, and $A_\pi(\boldsymbol{s}_t,\boldsymbol{a}_t)$ provides a measure of how each action compares to the average performance at the state $\boldsymbol{s}_t$, which is given by $V_\pi(\boldsymbol{s}_t)$. Using $A_\pi(\boldsymbol{s}_t,\boldsymbol{a}_t)$ centers the learning signal and reduces variance significantly.

Besides high variance, the main problem with policy gradient is that it requires on-policy samples which makes it very sample intensive. To achieve similar sample efficiency as off-policy methods, it is naturally crucial to use off-policy data, but it is nontrivial. Prior attempts use importance sampling to use off-policy trajectories; however, these are known to be difficult scale to high-dimensional action spaces because of rapidly degenerating importance weights (Precup, 2000).

2.2 Off-policy Actor-Critic Methods

Actor-critic methods (Sutton et al., 1999) include a *policy evaluation* step, which uses temporal difference (TD) learning to fit a critic Q_w for the current policy $\pi(\theta)$, and a *policy improvement* step which greedily optimizes the policy π against the critic estimate Q_w. Significant gain in sample efficiency is achievable using off-policy TD learning for the critic, as in Q-learning and deterministic policy gradient (Sutton, 1990; Silver et al., 2014), recently popularized by *experience replay* for training deep Q networks (Mnih et al., 2015; Lillicrap et al., 2016; Gu et al., 2016b).

Deep deterministic policy gradient (DDPG) (Silver et al., 2014; Lillicrap et al., 2016) is an instance of off-policy algorithms which achieves significant results on high-dimensional continuous control tasks. The updates for this method are given below, where $\pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t) = \delta(\boldsymbol{a}_t = \boldsymbol{\mu}_\theta(\boldsymbol{s}_t))$ is a deterministic policy, β is arbitrary exploration distribution, and ρ_β corresponds to sampling from a replay buffer:

$$\begin{aligned} w &= \arg\min_w \mathbb{E}_{\boldsymbol{s}_t\sim\rho_\beta(\cdot),\boldsymbol{a}_t\sim\beta(\cdot|\boldsymbol{s}_t)}[(r(\boldsymbol{s}_t,\boldsymbol{a}_t) + \gamma Q(\boldsymbol{s}_{t+1},\boldsymbol{\mu}_\theta(\boldsymbol{s}_{t+1})) - Q_w(\boldsymbol{s}_t,\boldsymbol{a}_t))^2] \\ \theta &= \arg\max_\theta \mathbb{E}_{\boldsymbol{s}_t\sim\rho_\beta(\cdot)}[Q_w(\boldsymbol{s}_t,\boldsymbol{\mu}_\theta(\boldsymbol{s}_t))]. \end{aligned} \tag{4}$$

When the critic and policy are parametrized with neural networks, full optimization is expensive, and instead stochastic gradient optimization is used. The gradient in the *policy improvement* phase is given below, which is generally a biased gradient of $J(\theta)$.

$$\nabla_\theta J(\theta) \approx \mathbb{E}_{\boldsymbol{s}_t\sim\rho_\beta(\cdot)}[\nabla_{\boldsymbol{a}} Q_w(\boldsymbol{s}_t,\boldsymbol{a})|_{\boldsymbol{a}=\boldsymbol{\mu}_\theta(\boldsymbol{s}_t)}\nabla_\theta \boldsymbol{\mu}_\theta(\boldsymbol{s}_t)] \tag{5}$$

The crucial benefits of DDPG are that it does not rely on high variance REINFORCE gradients and is trainable on off-policy data. These properties make DDPG and other analogous off-policy methods significantly more sample-efficient than policy gradient methods (Lillicrap et al., 2016; Gu et al., 2016b; Duan et al., 2016). However, the use of a biased policy gradient estimator makes analyzing its convergence and stability properties difficult.

3 Q-Prop

In this section, we derive the Q-Prop estimator for policy gradient. The key idea from this estimator comes from observing Equations 2 and 5 and noting that the former provides an unbiased, but

high variance gradient, while the latter provides a deterministic, but biased gradient. By using the deterministic biased estimator as a particular form of control variate (Ross, 2006; Paisley et al., 2012) for the unbiased policy gradient estimator, we can effectively use both types of gradient information to construct a new estimator that is in general unbiased, and in practice exhibits improved sample efficiency through the inclusion of off-policy samples.

3.1 Q-Prop Estimator

To derive the Q-Prop gradient estimator, we start by using the first-order Taylor expansion of an arbitrary function $f(\boldsymbol{s}_t, \boldsymbol{a}_t)$, $\bar{f}(\boldsymbol{s}_t, \boldsymbol{a}_t) = f(\boldsymbol{s}_t, \bar{\boldsymbol{a}}_t) + \nabla_{\boldsymbol{a}} f(\boldsymbol{s}_t, \boldsymbol{a})|_{\boldsymbol{a}=\bar{\boldsymbol{a}}_t}(\boldsymbol{a}_t - \bar{\boldsymbol{a}}_t)$, as the control variate for the policy gradient estimator. We use $\hat{Q}(\boldsymbol{s}_t, \boldsymbol{a}_t) = \sum_{t'=t}^{t} \gamma^{t'-t} r(\boldsymbol{s}_{t'}, \boldsymbol{a}_{t'})$ to denote Monte Carlo return from state $\boldsymbol{s}_t$ and action $\boldsymbol{a}_t$, i.e. $\mathbb{E}_\pi[\hat{Q}(\boldsymbol{s}_t, \boldsymbol{a}_t)] = r(\boldsymbol{s}_t, \boldsymbol{a}_t) + \gamma \mathbb{E}_p[V_\pi(\boldsymbol{s}_{t+1})]$, and $\boldsymbol{\mu}_\theta(\boldsymbol{s}_t) = \mathbb{E}_{\pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)}[\boldsymbol{a}_t]$ to denote the expected action of a stochastic policy π_θ. Full derivation is in Appendix A.

$$\begin{aligned} \nabla_\theta J(\theta) &= \mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)(\hat{Q}(\boldsymbol{s}_t, \boldsymbol{a}_t) - \bar{f}(\boldsymbol{s}_t, \boldsymbol{a}_t))] + \mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)\bar{f}(\boldsymbol{s}_t, \boldsymbol{a}_t)] \\ &= \mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)(\hat{Q}(\boldsymbol{s}_t, \boldsymbol{a}_t) - \bar{f}(\boldsymbol{s}_t, \boldsymbol{a}_t))] + \mathbb{E}_{\rho_\pi}[\nabla_{\boldsymbol{a}} f(\boldsymbol{s}_t, \boldsymbol{a})|_{\boldsymbol{a}=\bar{\boldsymbol{a}}_t} \nabla_\theta \boldsymbol{\mu}_\theta(\boldsymbol{s}_t)] \end{aligned} \tag{6}$$

Eq. 6 is general for arbitrary function $f(\boldsymbol{s}_t, \boldsymbol{a}_t)$ that is differentiable with respect to $\boldsymbol{a}_t$ at an arbitrary value of $\bar{\boldsymbol{a}}_t$; however, a sensible choice is to use the critic Q_w for f and $\boldsymbol{\mu}_\theta(\boldsymbol{s}_t)$ for $\bar{\boldsymbol{a}}_t$ to get,

$$\nabla_\theta J(\theta) = \mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)(\hat{Q}(\boldsymbol{s}_t, \boldsymbol{a}_t) - \bar{Q}_w(\boldsymbol{s}_t, \boldsymbol{a}_t))] + \mathbb{E}_{\rho_\pi}[\nabla_{\boldsymbol{a}} Q_w(\boldsymbol{s}_t, \boldsymbol{a})|_{\boldsymbol{a}=\boldsymbol{\mu}_\theta(\boldsymbol{s}_t)} \nabla_\theta \boldsymbol{\mu}_\theta(\boldsymbol{s}_t)]. \tag{7}$$

Finally, since in practice we estimate advantages $\hat{A}(\boldsymbol{s}_t, \boldsymbol{a}_t)$, we write the Q-Prop estimator in terms of advantages to complete the basic derivation,

$$\begin{aligned} &\nabla_\theta J(\theta) = \mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)(\hat{A}(\boldsymbol{s}_t, \boldsymbol{a}_t) - \bar{A}_w(\boldsymbol{s}_t, \boldsymbol{a}_t))] + \mathbb{E}_{\rho_\pi}[\nabla_{\boldsymbol{a}} Q_w(\boldsymbol{s}_t, \boldsymbol{a})|_{\boldsymbol{a}=\boldsymbol{\mu}_\theta(\boldsymbol{s}_t)} \nabla_\theta \boldsymbol{\mu}_\theta(\boldsymbol{s}_t)] \\ &\bar{A}(\boldsymbol{s}_t, \boldsymbol{a}_t) = \bar{Q}(\boldsymbol{s}_t, \boldsymbol{a}_t) - \mathbb{E}_{\pi_\theta}[\bar{Q}(\boldsymbol{s}_t, \boldsymbol{a}_t)] = \nabla_{\boldsymbol{a}} Q_w(\boldsymbol{s}_t, \boldsymbol{a})|_{\boldsymbol{a}=\boldsymbol{\mu}_\theta(\boldsymbol{s}_t)}(\boldsymbol{a}_t - \boldsymbol{\mu}_\theta(\boldsymbol{s}_t)). \end{aligned} \tag{8}$$

Eq. 8 comprises of an analytic gradient through the critic as in Eq. 5 and a residual REINFORCE gradient in Eq. 2. From the above derivation, Q-Prop is simply an unbiased policy gradient estimator with a special form of control variate. The important insight comes from the fact that Q_w can be trained using off-policy data as in Eq. 4. Under this setting, Q-Prop is no longer just a policy gradient method, but more closely resembles an actor-critic method, except in *policy improvement* step it has an additional REINFORCE correction term that ensures that the gradient estimator is unbiased regardless of the parametrization, training method, and performance of the critic.

Intuitively, if the critic Q_w approximates Q_π well, it provides a reliable gradient, reduces the estimator variance, and improves the convergence rate. Interestingly, control variate analysis in the next section shows that this is not the only circumstance where Q-Prop helps reduce variance.

3.2 Control Variate Analysis and Adaptive Q-Prop

For Q-Prop to be applied reliably, it is crucial to analyze how the variance of the estimator changes before and after the application of control variate. Following the prior work on control variate (Ross, 2006; Paisley et al., 2012), we first introduce $\eta(\boldsymbol{s}_t)$ to Eq. 8, a weighing variable that modulates the strength of control variate. This additional variable $\eta(\boldsymbol{s}_t)$ does not introduce bias to the estimator.

$$\begin{aligned} \nabla_\theta J(\theta) =& \mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)(\hat{A}(\boldsymbol{s}_t, \boldsymbol{a}_t) - \eta(\boldsymbol{s}_t)\bar{A}_w(\boldsymbol{s}_t, \boldsymbol{a}_t))] \\ &+ \mathbb{E}_{\rho_\pi}[\eta(\boldsymbol{s}_t)\nabla_{\boldsymbol{a}} Q_w(\boldsymbol{s}_t, \boldsymbol{a})|_{\boldsymbol{a}=\boldsymbol{\mu}_\theta(\boldsymbol{s}_t)} \nabla_\theta \boldsymbol{\mu}_\theta(\boldsymbol{s}_t)] \end{aligned} \tag{9}$$

A measure of the variance of this estimator is below, where $m = 1...M$ indexes the dimension of θ,

$$\text{Var}^* = \mathbb{E}_{\rho_\pi}[\sum_m \text{Var}_{\boldsymbol{a}_t}(\nabla_{\theta_m} \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)(\hat{A}(\boldsymbol{s}_t, \boldsymbol{a}_t) - \eta(\boldsymbol{s}_t)\bar{A}(\boldsymbol{s}_t, \boldsymbol{a}_t)))]. \tag{10}$$

If we choose $\eta(\boldsymbol{s}_t)$ such that $\text{Var}^* < \text{Var}$, where $\text{Var} = \mathbb{E}_{\rho_\pi}[\sum_m \text{Var}_{\boldsymbol{a}_t}(\nabla_{\theta_m} \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)\hat{A}(\boldsymbol{s}_t, \boldsymbol{a}_t))]$ is the original estimator variance measure, then we have managed to reduce the variance. Directly analyzing the above variance measures is nontrivial, just like computing the optimal baseline is

difficult (Williams, 1992). In addition, it is often impractical get multiple action samples from the same state, which prohibits using naïve Monte Carlo to estimate the expectations. Thus, we propose a surrogate variance measure, $\text{Var} = \mathbb{E}_{\rho_\pi}[\text{Var}_{\boldsymbol{a}_t}(\hat{A}(\boldsymbol{s}_t, \boldsymbol{a}_t))]$. A similar surrogate is also used by prior work on learning state-dependent baseline (Mnih & Gregor, 2014), and the benefit is that the measure becomes more tractable,

$$\begin{aligned} \text{Var}^* &= \mathbb{E}_{\rho_\pi}[\text{Var}_{\boldsymbol{a}_t}(\hat{A}(\boldsymbol{s}_t, \boldsymbol{a}_t) - \eta(\boldsymbol{s}_t)\bar{A}(\boldsymbol{s}_t, \boldsymbol{a}_t))] \\ &= \text{Var} + \mathbb{E}_{\rho_\pi}[-2\eta(\boldsymbol{s}_t)\text{Cov}_{\boldsymbol{a}_t}(\hat{A}(\boldsymbol{s}_t, \boldsymbol{a}_t), \bar{A}(\boldsymbol{s}_t, \boldsymbol{a}_t)) + \eta(\boldsymbol{s}_t)^2 \text{Var}_{\boldsymbol{a}_t}(\bar{A}(\boldsymbol{s}_t, \boldsymbol{a}_t))]. \end{aligned} \tag{11}$$

Since $\mathbb{E}_\pi[\hat{A}(\boldsymbol{s}_t, \boldsymbol{a}_t)] = \mathbb{E}_\pi[\bar{A}(\boldsymbol{s}_t, \boldsymbol{a}_t)] = 0$, the terms can be simplified as below,

$$\begin{aligned} \text{Cov}_{\boldsymbol{a}_t}(\hat{A}, \bar{A}) &= \mathbb{E}_\pi[\hat{A}(\boldsymbol{s}_t, \boldsymbol{a}_t)\bar{A}(\boldsymbol{s}_t, \boldsymbol{a}_t)] \\ \text{Var}_{\boldsymbol{a}_t}(\bar{A}) &= \mathbb{E}_\pi[\bar{A}(\boldsymbol{s}_t, \boldsymbol{a}_t)^2] = \nabla_{\boldsymbol{a}} Q_w(\boldsymbol{s}_t, \boldsymbol{a})|^T_{\boldsymbol{a}=\boldsymbol{\mu}_\theta(\boldsymbol{s}_t)} \Sigma_\theta(\boldsymbol{s}_t) \nabla_{\boldsymbol{a}} Q_w(\boldsymbol{s}_t, \boldsymbol{a})|_{\boldsymbol{a}=\boldsymbol{\mu}_\theta(\boldsymbol{s}_t)}, \end{aligned} \tag{12}$$

where $\Sigma_\theta(\boldsymbol{s}_t)$ is the covariance matrix of the stochastic policy π_θ. The nice property of Eq. 11 is that $\text{Var}_{\boldsymbol{a}_t}(\bar{A})$ is analytical and $\text{Cov}_{\boldsymbol{a}_t}(\hat{A}, \bar{A})$ can be estimated with single action sample. Using this estimate, we propose adaptive variants of Q-Prop that regulate the variance of the gradient estimate.

Adaptive Q-Prop. The optimal state-dependent factor $\eta(\boldsymbol{s}_t)$ can be computed per state, according to $\eta^*(\boldsymbol{s}_t) = \text{Cov}_{\boldsymbol{a}_t}(\hat{A}, \bar{A})/\text{Var}_{\boldsymbol{a}_t}(\bar{A})$. This provides maximum reduction in variance according to Eq. 11. Substituting $\eta^*(\boldsymbol{s}_t)$ into Eq. 11, we get $\text{Var}^* = \mathbb{E}_{\rho_\pi}[(1 - \rho_{corr}(\hat{A}, \bar{A})^2)\text{Var}_{\boldsymbol{a}_t}(\hat{A})]$, where ρ_{corr} is the correlation coefficient, which achieves guaranteed variance reduction if at any state $\bar{A}$ is correlated with $\hat{A}$. We call this the fully adaptive Q-Prop method. An important conclusion from this analysis is that, in adaptive Q-Prop, the critic Q_w does not necessarily need to be approximating Q_π well to produce good results. Its Taylor expansion merely needs to be correlated with $\hat{A}$, positively or even negatively. This is in contrast with actor-critic methods, where performance is greatly dependent on the absolute accuracy of the critic's approximation.

Conservative and Aggressive Q-Prop. In practice, the single-sample estimate of $\text{Cov}_{\boldsymbol{a}_t}(\hat{A}, \bar{A})$ has high variance itself, and we propose the following two practical implementations of adaptive Q-Prop: (1) $\eta(\boldsymbol{s}_t) = 1$ if $\hat{\text{Cov}}_{\boldsymbol{a}_t}(\hat{A}, \bar{A}) > 0$ and $\eta(\boldsymbol{s}_t) = 0$ if otherwise, and (2) $\eta(\boldsymbol{s}_t) = \text{sign}(\hat{\text{Cov}}_{\boldsymbol{a}_t}(\hat{A}, \bar{A}))$. The first implementation, which we call conservative Q-Prop, can be thought of as a more conservative version of Q-Prop, which effectively disables the control variate for some samples of the states. This is sensible as if $\hat{A}$ and $\bar{A}$ are negatively correlated, it is likely that the critic is very poor. The second variant can correspondingly be termed aggressive Q-Prop, since it makes more liberal use of the control variate.

Algorithm 1 Adaptive Q-Prop

Initialize w for critic Q_w, θ for stochastic policy π_θ, and replay buffer $\mathscr{R} \leftarrow \emptyset$.
repeat
  for $e = 1, \ldots, E$ **do** ▷ Collect E episodes of on-policy experience using π_θ
    $\boldsymbol{s}_{0,e} \sim p(\boldsymbol{s}_0)$
    for $t = 0, \ldots, T-1$ **do**
      $\boldsymbol{a}_{t,e} \sim \pi_\theta(\cdot|\boldsymbol{s}_{t,e})$, $\boldsymbol{s}_{t+1,e} \sim p(\cdot|\boldsymbol{s}_{t,e}, \boldsymbol{a}_{t,e})$, $r_{t,e} = r(\boldsymbol{s}_{t,e}, \boldsymbol{a}_{t,e})$
  Add batch data $\mathscr{B} = \{\boldsymbol{s}_{0:T,1:E}, \boldsymbol{a}_{0:T-1,1:E}, r_{0:T-1,1:E}\}$ to replay buffer $\mathscr{R}$
  Take $E \cdot T$ gradient steps on Q_w using $\mathscr{R}$ and π_θ
  Fit $V_\phi(\boldsymbol{s}_t)$ using $\mathscr{B}$
  Compute $\hat{A}_{t,e}$ using GAE(λ) and $\bar{A}_{t,e}$ using Eq. 7
  Set $\eta_{t,e}$ based on Section 3.2
  Compute and center the learning signals $l_{t,e} = \hat{A}_{t,e} - \eta_{t,e}\bar{A}_{t,e}$
  Compute $\nabla_\theta J(\theta) \approx \frac{1}{ET}\sum_e \sum_t \nabla_\theta \log \pi_\theta(\boldsymbol{a}_{t,e}|\boldsymbol{s}_{t,e}) l_{t,e} + \eta_{t,e} \nabla_{\boldsymbol{a}} Q_w(\boldsymbol{s}_{t,e}, \boldsymbol{a})|_{\boldsymbol{a}=\boldsymbol{\mu}_\theta(\boldsymbol{s}_{t,e})} \nabla_\theta \boldsymbol{\mu}_\theta(\boldsymbol{s}_{t,e})$
  Take a gradient step on π_θ using $\nabla_\theta J(\theta)$, optionally with a trust-region constraint using $\mathscr{B}$
until π_θ converges.

3.3 Q-PROP ALGORITHM

Pseudo-code for the adaptive Q-Prop algorithm is provided in Algorithm 1. It is a mixture of policy gradient and actor-critic. At each iteration, it first rolls out the stochastic policy to collect on-policy

samples, adds the batch to a replay buffer, takes a few gradient steps on the critic, computes $\hat{A}$ and $\bar{A}$, and finally applies a gradient step on the policy π_θ. In our implementation, the critic Q_w is fitted with off-policy TD learning using the same techniques as in DDPG (Lillicrap et al., 2016):

$$w = \arg\min_w \mathbb{E}_{\boldsymbol{s}_t \sim \rho_\beta(\cdot), \boldsymbol{a}_t \sim \beta(\cdot|\boldsymbol{s}_t)}[(r(\boldsymbol{s}_t, \boldsymbol{a}_t) + \gamma \mathbb{E}_\pi[Q'(\boldsymbol{s}_{t+1}, \boldsymbol{a}_{t+1})] - Q_w(\boldsymbol{s}_t, \boldsymbol{a}_t))^2]. \tag{13}$$

V_ϕ is fitted with the same technique in (Schulman et al., 2016). Generalized advantage estimation (GAE) (Schulman et al., 2016) is used to estimate $\hat{A}$. The policy update can be done by any method that utilizes the first-order gradient and possibly the on-policy batch data, which includes trust region policy optimization (TRPO) (Schulman et al., 2015). Importantly, this is just one possible implementation of Q-Prop, and in Appendix C we show a more general form that can interpolate between pure policy gradient and off-policy actor-critic.

A limitation with Algorithm 1 is that computation time is significantly more than TRPO because it requires many gradient steps to train Q_w at each iteration. Fortunately, training routines for Q_w can be made asynchronous. Furthermore, in real-world applications, data collection speed is often the bottleneck, enabling sufficient time between policy iterations to fit Q_w well with the updated policy π_θ and the replay data.

4 Related Work

Variance reduction in policy gradient methods is a long-standing problem with a large body of prior work (Williams, 1992; Greensmith et al., 2004; Schulman et al., 2016). However, exploration of action-dependent control variates is relatively recent, with most work focusing instead on simpler baselining techniques (Ross, 2006). A subtle exception is compatible feature approximation (Sutton et al., 1999) which can be viewed as a control variate as explained in Appendix B. Another exception is doubly robust estimator in contextual bandits (Dudík et al., 2011), which uses a different control variate whose bias cannot be tractably corrected. Control variates were explored recently not in RL but for approximate inference in stochastic models (Paisley et al., 2012), and the closest related work in that domain is the MuProp algorithm (Gu et al., 2016a) which uses a mean-field network as a surrogate for backpropagating a deterministic gradient through stochastic discrete variables. MuProp is not directly applicable to model-free RL because the dynamics are unknown; however, it can be if the dynamics are learned as in model-based RL (Atkeson & Santamaria, 1997; Deisenroth & Rasmussen, 2011). This model-based Q-Prop is itself an interesting direction of research as it effectively corrects bias in model-based learning.

Part of the benefit of Q-Prop is the ability to use off-policy data to improve on-policy policy gradient methods. Prior methods that combine off-policy data with policy gradients either introduce bias (Sutton et al., 1999; Silver et al., 2014) or use importance weighting, which is known to result in degenerate importance weights in high dimensions, resulting in very high variance (Precup, 2000; Levine & Koltun, 2013). Q-Prop provides a new approach for using off-policy data to reduce variance, while remaining unbiased.

5 Experiments

(a) (b) (c) (d) (e) (f) (g)

Figure 1: Illustrations of OpenAI Gym MuJoCo domains (Brockman et al., 2016; Duan et al., 2016): (a) Ant, (b) HalfCheetah, (c) Hopper, (d) Humanoid, (e) Reacher, (f) Swimmer, (g) Walker.

We evaluated Q-Prop and its variants on continuous control environments from the OpenAI Gym benchmark (Brockman et al., 2016) using the MuJoCo physics simulator (Todorov et al., 2012) as shown in Figure 1. Algorithms are identified by acronyms, followed by a number indicating batch

size, except for DDPG, which is a prior online actor-critic algorithm (Lillicrap et al., 2016). "c-" and "v-" denote conservative and aggressive Q-Prop variants as described in Section 3.2. "TR-" denotes trust-region policy optimization (Schulman et al., 2015), while "V-" denotes vanilla policy gradient. For example, "TR-c-Q-Prop-5000" means convervative Q-Prop with the trust-region policy update, and a batch size of 5000. "VPG" and "TRPO" are vanilla policy gradient and trust-region policy optimization respectively (Schulman et al., 2016; Duan et al., 2016). Unless otherwise stated, all policy gradient methods are implemented with GAE($\lambda = 0.97$) (Schulman et al., 2016). Note that TRPO-GAE is currently the state-of-the-art method on most of the OpenAI Gym benchmark tasks, though our experiments show that a well-tuned DDPG implementation sometimes achieves better results. Our algorithm implementations are built on top of the rllab TRPO and DDPG codes from Duan et al. (2016) and will be released upon publication. Policy and value function architectures and other training details including hyperparameter values are provided in Appendix D.

5.1 Adaptive Q-Prop

First, it is useful to identify how reliable each variant of Q-Prop is. In this section, we analyze standard Q-Prop and two adaptive variants, c-Q-Prop and a-Q-Prop, and demonstrate the stability of the method across different batch sizes. Figure 2a shows a comparison of Q-Prop variants with trust-region updates on the HalfCheetah-v1 domain, along with the best performing TRPO hyperparameters. The results are consistent with theory: conservative Q-Prop achieves much more stable performance than the standard and aggressive variants, and all Q-Prop variants significantly outperform TRPO in terms of sample efficiency, e.g. conservative Q-Prop reaches average reward of 4000 using about 10 times less samples than TRPO.

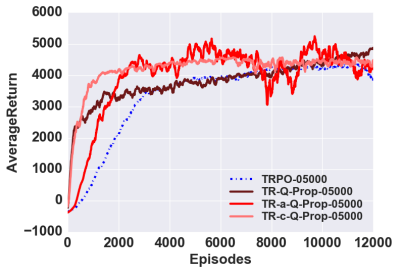

(a) Standard Q-Prop vs adaptive variants.

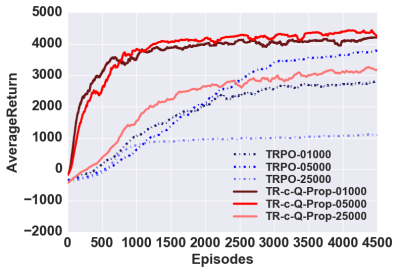

(b) Conservative Q-Prop vs TRPO across batch sizes.

Figure 2: Average return over episodes in HalfCheetah-v1 during learning, exploring adaptive Q-Prop methods and different batch sizes. All variants of Q-Prop substantially outperform TRPO in terms of sample efficiency. TR-c-QP, conservative Q-Prop with trust-region update performs most stably across different batch sizes.

Figure 2b shows the performance of conservative Q-Prop against TRPO across different batch sizes. Due to high variance in gradient estimates, TRPO typically requires very large batch sizes, e.g. 25000 steps or 25 episodes per update, to perform well. We show that our Q-Prop methods can learn even with just 1 episode per update, and achieves better sample efficiency with small batch sizes. This shows that Q-Prop significantly reduces the variance compared to the prior methods.

As we discussed in Section 1, stability is a significant challenge with state-of-the-art deep RL methods, and is very important for being able to reliably use deep RL for real world tasks. In the rest of the experiments, we will use conservative Q-Prop as the main Q-Prop implementation.

5.2 Evaluation Across Algorithms

In this section, we evaluate two versions of conservative Q-Prop, v-c-Q-Prop using vanilla policy gradient and TR-c-Q-Prop using trust-region updates, against other model-free algorithms on the HalfCheetah-v1 domain. Figure 3a shows that c-Q-Prop methods significantly outperform the best TRPO and VPG methods. Even Q-Prop with vanilla policy gradient is comparable to TRPO, confirming the significant benefits from variance reduction. DDPG on the other hand exhibits inconsistent performances. With proper reward scaling, i.e. "DDPG-r0.1", it outperforms other methods

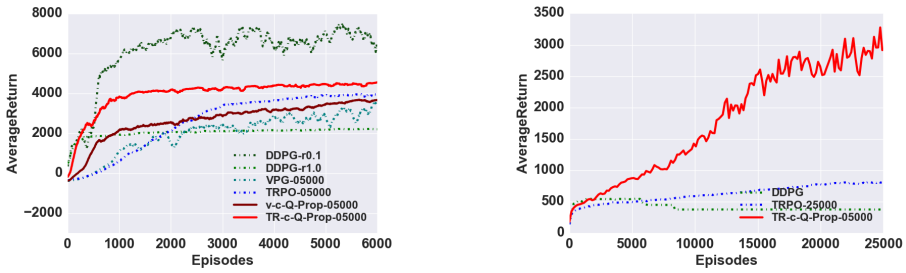

(a) Comparing algorithms on HalfCheetah-v1.

(b) Comparing algorithms on Humanoid-v1.

Figure 3: Average return over episodes in HalfCheetah-v1 and Humanoid-v1 during learning, comparing Q-Prop against other model-free algorithms. Q-Prop with vanilla policy gradient outperforms TRPO on HalfCheetah. Q-Prop significantly outperforms TRPO in convergence time on Humanoid.

as well as the DDPG results reported in prior work (Duan et al., 2016; Amos et al., 2016). This illustrates the sensitivity of DDPG to hyperparameter settings, while Q-Prop exhibits more stable, monotonic learning behaviors when compared to DDPG. In the next section we show this improved stability allows Q-Prop to outperform DDPG in more complex domains.

5.3 Evaluation across domains

Lastly, we evaluate Q-Prop against TRPO and DDPG across multiple domains. While the gym environments are biased toward locomotion, we expect we can achieve similar performance on manipulation tasks such as those in Lillicrap et al. (2016). Table 1 summarizes the results, including the best attained average rewards and the steps to convergence. Q-Prop consistently outperform TRPO in terms of sample complexity and sometimes achieves higher rewards than DDPG in more complex domains. A particularly notable case is shown in Figure 3b, where Q-Prop substantially improves sample efficiency over TRPO on Humanoid-v1 domain, while DDPG cannot find a good solution.

The better performance on the more complex domains highlights the importance of stable deep RL algorithms: while costly hyperparameter sweeps may allow even less stable algorithms to perform well on simpler problems, more complex tasks might have such narrow regions of stable hyperparameters that discovering them becomes impractical.

		TR-c-Q-Prop		TRPO		DDPG	
Domain	Threshold	MaxReturn.	Episodes	MaxReturn	Epsisodes	MaxReturn	Episodes
Ant	3500	3534	**4975**	**4239**	13825	957	N/A
HalfCheetah	4700	4811	20785	4734	26370	**7490**	**600**
Hopper	2000	**2957**	5945	2486	5715	2604	**965**
Humanoid	2500	**>3492**	**14750**	918	>30000	552	N/A
Reacher	-7	**-6.0**	**2060**	-6.7	2840	-6.6	**1800**
Swimmer	90	103	2045	110	3025	**150**	**500**
Walker	3000	**4030**	3685	3567	18875	3626	**2125**

Table 1: Q-Prop, TRPO and DDPG results showing the max average rewards attained in the first 30k episodes and the episodes to cross specific reward thresholds. Q-Prop often learns more sample efficiently than TRPO and can solve difficult domains such as Humanoid better than DDPG.

6 Discussion and Conclusion

We presented Q-Prop, a policy gradient algorithm that combines reliable, consistent, and potentially unbiased on-policy gradient estimation with a sample-efficient off-policy critic that acts as a control variate. The method provides a large improvement in sample efficiency compared to state-of-the-art policy gradient methods such as TRPO, while outperforming state-of-the-art actor-critic methods on more challenging tasks such as humanoid locomotion. We hope that techniques like these, which combine unbiased gradient estimation with sample-efficient variance reduction through off-policy critics, will eventually lead to deep reinforcement learning algorithms that are more stable and efficient, and therefore better suited for application to complex real-world learning tasks.

ACKNOWLEDGMENTS

We thank Rocky Duan for answering questions about rllab code, and Yutian Chen and Laurent Dinh for discussion on control variates. SG was funded by NSERC and a Google Focused Research Award.

REFERENCES

Brandon Amos, Lei Xu, and J Zico Kolter. Input convex neural networks. *arXiv preprint arXiv:1609.07152*, 2016.

Christopher G Atkeson and Juan Carlos Santamaria. A comparison of direct and model-based reinforcement learning. In *In International Conference on Robotics and Automation*. Citeseer, 1997.

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

Marc Deisenroth and Carl E Rasmussen. Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pp. 465–472, 2011.

Yan Duan, Xi Chen, Rein Houthooft, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. *International Conference on Machine Learning (ICML)*, 2016.

Miroslav Dudík, John Langford, and Lihong Li. Doubly robust policy evaluation and learning. *arXiv preprint arXiv:1103.4601*, 2011.

Evan Greensmith, Peter L Bartlett, and Jonathan Baxter. Variance reduction techniques for gradient estimates in reinforcement learning. *Journal of Machine Learning Research*, 5(Nov):1471–1530, 2004.

Shixiang Gu, Sergey Levine, Ilya Sutskever, and Andriy Mnih. Muprop: Unbiased backpropagation for stochastic neural networks. *International Conference on Learning Representations (ICLR)*, 2016a.

Shixiang Gu, Tim Lillicrap, Ilya Sutskever, and Sergey Levine. Continuous deep q-learning with model-based acceleration. In *International Conference on Machine Learning (ICML)*, 2016b.

Hado V Hasselt. Double q-learning. In *Advances in Neural Information Processing Systems*, pp. 2613–2621, 2010.

Sham Kakade. A natural policy gradient. In *NIPS*, volume 14, pp. 1531–1538, 2001.

Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Guy Lever. Deterministic policy gradient algorithms. 2014.

Sergey Levine and Vladlen Koltun. Guided policy search. In *International Conference on Machine Learning (ICML)*, pp. 1–9, 2013.

Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *International Conference on Learning Representations (ICLR)*, 2016.

Andriy Mnih and Karol Gregor. Neural variational inference and learning in belief networks. *International Conference on Machine Learning (ICML)*, 2014.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy P Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2016.

Rémi Munos, Tom Stepleton, Anna Harutyunyan, and Marc G Bellemare. Safe and efficient off-policy reinforcement learning. *arXiv preprint arXiv:1606.02647*, 2016.

John Paisley, David Blei, and Michael Jordan. Variational bayesian inference with stochastic search. *International Conference on Machine Learning (ICML)*, 2012.

Jan Peters and Stefan Schaal. Policy gradient methods for robotics. In *International Conference on Intelligent Robots and Systems (IROS)*, pp. 2219–2225. IEEE, 2006.

Jan Peters, Katharina Mülling, and Yasemin Altun. Relative entropy policy search. In *AAAI*. Atlanta, 2010.

Doina Precup. Eligibility traces for off-policy policy evaluation. *Computer Science Department Faculty Publication Series*, pp. 80, 2000.

Sheldon M Ross. Simulation. *Burlington, MA: Elsevier*, 2006.

John Schulman, Sergey Levine, Pieter Abbeel, Michael I. Jordan, and Philipp Moritz. Trust region policy optimization. In *International Conference on Machine Learning (ICML)*, pp. 1889–1897, 2015.

John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *International Conference on Learning Representations (ICLR)*, 2016.

David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *International Conference on Machine Learning (ICML)*, 2014.

David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

Richard S Sutton. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *International Conference on Machine Learning (ICML)*, pp. 216–224, 1990.

Richard S Sutton, David A McAllester, Satinder P Singh, Yishay Mansour, et al. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems (NIPS)*, volume 99, pp. 1057–1063, 1999.

Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033. IEEE, 2012.

Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.

Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.

A Q-PROP ESTIMATOR DERIVATION

The full derivation of the Q-Prop estimator is shown in Eq. 14. We make use of the following property that is commonly used in baseline derivations:

$$\mathbb{E}_{p_\theta(x)}[\nabla_\theta \log p_\theta(x)] = \int_x \nabla_\theta p_\theta(x) = \nabla_\theta \int_x p(x) = 0$$

This holds true when $f(\boldsymbol{s}_t, \boldsymbol{a}_t)$ is an arbitrary function differentiable with respect to $\boldsymbol{a}_t$ and $\bar{f}$ is its first-order Taylor expansion around $\boldsymbol{a}_t = \bar{\boldsymbol{a}}_t$, i.e. $\bar{f}(\boldsymbol{s}_t, \boldsymbol{a}_t) = f(\boldsymbol{s}_t, \bar{\boldsymbol{a}}_t) + \nabla_{\boldsymbol{a}} f(\boldsymbol{s}_t, \boldsymbol{a})|_{\boldsymbol{a}=\bar{\boldsymbol{a}}_t}(\boldsymbol{a}_t - \bar{\boldsymbol{a}}_t)$. Here, $\boldsymbol{\mu}_\theta(\boldsymbol{s}_t) = \mathbb{E}_\pi[\boldsymbol{a}_t]$ is the mean of stochastic policy π_θ. The derivation appears below:

$$
\begin{aligned}
\nabla_\theta J(\theta) &= \mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)(\hat{Q}(\boldsymbol{s}_t,\boldsymbol{a}_t) - \bar{f}(\boldsymbol{s}_t,\boldsymbol{a}_t)] + \mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)\bar{f}(\boldsymbol{s}_t,\boldsymbol{a}_t)] \\
g(\theta) &= \mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)\bar{f}(\boldsymbol{s}_t,\boldsymbol{a}_t)] \\
&= \mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)(f(\boldsymbol{s}_t,\bar{\boldsymbol{a}}_t) + \nabla_{\boldsymbol{a}} f(\boldsymbol{s}_t,\boldsymbol{a})|_{\boldsymbol{a}=\bar{\boldsymbol{a}}_t}(\boldsymbol{a}_t - \bar{\boldsymbol{a}}_t))] \\
&= \mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)\nabla_{\boldsymbol{a}} f(\boldsymbol{s}_t,\boldsymbol{a})|_{\boldsymbol{a}=\bar{\boldsymbol{a}}_t}\boldsymbol{a}_t] \\
&= \mathbb{E}_{\rho_\pi}\left[\int_{\boldsymbol{a}_t} \nabla_\theta \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)\nabla_{\boldsymbol{a}} f(\boldsymbol{s}_t,\boldsymbol{a})|_{\boldsymbol{a}=\bar{\boldsymbol{a}}_t}\boldsymbol{a}_t\right] \\
&= \mathbb{E}_{\rho_\pi}\left[\nabla_{\boldsymbol{a}} f(\boldsymbol{s}_t,\boldsymbol{a})|_{\boldsymbol{a}=\bar{\boldsymbol{a}}_t}\int_{\boldsymbol{a}_t} \nabla_\theta \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)\boldsymbol{a}_t\right] \\
&= \mathbb{E}_{\rho_\pi}[\nabla_{\boldsymbol{a}} f(\boldsymbol{s}_t,\boldsymbol{a})|_{\boldsymbol{a}=\bar{\boldsymbol{a}}_t}\nabla_\theta \mathbb{E}_\pi[\boldsymbol{a}_t]] \\
&= \mathbb{E}_{\rho_\pi}[\nabla_{\boldsymbol{a}} f(\boldsymbol{s}_t,\boldsymbol{a})|_{\boldsymbol{a}=\bar{\boldsymbol{a}}_t}\nabla_\theta \boldsymbol{\mu}_\theta(\boldsymbol{s}_t)] \\
\nabla_\theta J(\theta) &= \mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)(\hat{Q}(\boldsymbol{s}_t,\boldsymbol{a}_t) - \bar{f}(\boldsymbol{s}_t,\boldsymbol{a}_t)] + g(\theta) \\
&= \mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)(\hat{Q}(\boldsymbol{s}_t,\boldsymbol{a}_t) - \bar{f}(\boldsymbol{s}_t,\boldsymbol{a}_t)] + \mathbb{E}_{\rho_\pi}[\nabla_{\boldsymbol{a}} f(\boldsymbol{s}_t,\boldsymbol{a})|_{\boldsymbol{a}=\bar{\boldsymbol{a}}_t}\nabla_\theta \boldsymbol{\mu}_\theta(\boldsymbol{s}_t)]
\end{aligned}
\tag{14}
$$

B Connection Between Q-Prop and Compatible Feature Approximation

In this section we show that actor-critic with compatible feature approximation is a form of control variate. A critic Q_w is *compatible* (Sutton et al., 1999) if it satisfies (1) $Q_w(\boldsymbol{s}_t, \boldsymbol{a}_t) = w^T \nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)$, i.e. $\nabla_w Q_w(\boldsymbol{s}_t, \boldsymbol{a}_t) = \nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)$, and (2) w is fit with objective $w = \arg\min_w L(w) = \arg\min_w \mathbb{E}_{\rho_\pi,\pi}[(\hat{Q}(\boldsymbol{s}_t,\boldsymbol{a}_t) - Q_w(\boldsymbol{s}_t,\boldsymbol{a}_t))^2]$, that is fitting Q_w on on-policy Monte Carlo returns. Condition (2) implies the following identity,

$$
\nabla_w L = 2\mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)(\hat{Q}(\boldsymbol{s}_t,\boldsymbol{a}_t) - Q_w(\boldsymbol{s}_t,\boldsymbol{a}_t))] = 0. \tag{15}
$$

In compatible feature approximation, it directly uses Q_w as control variate, rather than its Taylor expansion $\bar{Q}_w$ as in Q-Prop. Using Eq. 15, the unbiased policy gradient is,

$$
\begin{aligned}
\nabla_\theta J(\theta) &= \mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t) Q_w(\boldsymbol{s}_t,\boldsymbol{a}_t)] \\
&= \mathbb{E}_{\rho_\pi,\pi}[(\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)^T)w] \\
&= \mathbb{E}_{\rho_\pi}[I(\theta;\boldsymbol{s}_t)w],
\end{aligned}
\tag{16}
$$

where $I(\theta;\boldsymbol{s}_t) = \mathbb{E}_{\pi_\theta}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)^T]$ is Fisher's information matrix. Thus, variance reduction depends on ability to compute or estimate $I(\theta;\boldsymbol{s}_t)$ and w effectively.

C Unifying Policy Gradient and Actor-Critic

Q-Prop closely ties together policy gradient and actor-critic algorithms. To analyze this point, we write a generalization of Eq. 9 below, introducing two additional variables α, ρ_{CR}:

$$
\begin{aligned}
\nabla_\theta J(\theta) \propto \alpha \mathbb{E}_{\rho_\pi,\pi}[\nabla_\theta \log \pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t)(\hat{A}(\boldsymbol{s}_t,\boldsymbol{a}_t) - \eta \bar{A}_w(\boldsymbol{s}_t,\boldsymbol{a}_t)] \\
+ \eta \mathbb{E}_{\rho_{CR}}[\nabla_{\boldsymbol{a}} Q_w(\boldsymbol{s}_t,\boldsymbol{a})|_{\boldsymbol{a}=\boldsymbol{\mu}_\theta(\boldsymbol{s}_t)}\nabla_\theta \boldsymbol{\mu}_\theta(\boldsymbol{s}_t)]
\end{aligned}
\tag{17}
$$

Eq. 17 enables more analysis where bias generally is introduced only when $\alpha \neq 1$ or $\rho_{CR} \neq \rho_\pi$. Importantly, Eq. 17 covers both policy gradient and deterministic actor-critic algorithm as its special cases. Standard policy gradient is recovered by $\eta = 0$, and deterministic actor-critic is recovered by $\alpha = 0$ and $\rho_{CR} = \rho_\beta$. This allows heuristic or automatic methods for dynamically changing these variables through the learning process for optimizing different metrics, e.g. sample efficiency, convergence speed, stability.

Table 2 summarizes the various edge cases of Eq. 17. For example, since we derive our method from a control variates standpoint, Q_w can be any function and the gradient remains unbiased. A natural

Parameter	Implementation options	Introduce bias?
Q_w	off-policy TD; on-policy TD(λ); model-based; etc.	No
V_ϕ	on-policy Monte Carlo fitting; $\mathbb{E}_{\pi_\theta}[Q_w(\boldsymbol{s}_t, \boldsymbol{a}_t)]$; etc	No
λ	$0 \leq \lambda \leq 1$	Yes, except $\lambda = 1$
α	$\alpha \geq 0$	Yes, except $\alpha = 1$
η	any η	No
ρ_{CR}	ρ of any policy	Yes, except $\rho_{CR} = \rho_\pi$

Table 2: Implementation options and edge cases of the generalized Q-Prop estimator in Eq. 17.

choice is to use off-policy temporal difference learning to learn the critic Q_w corresponding to policy π. This enables effectively utilizing off-policy samples without introducing bias. An interesting alternative to this is to utilize model-based roll-outs to estimate the critic, which resembles MuProp in stochastic neural networks (Gu et al., 2016a). Unlike prior work on using fitted dynamics model to accelerate model-free learning (Gu et al., 2016b), this approach does not introduce bias to the gradient of the original objective.

D Experiment Details

Policy and value function architectures. The network architectures are largely based on the benchmark paper by Duan et al. (2016). For policy gradient methods, the stochastic policy $\pi_\theta(\boldsymbol{a}_t|\boldsymbol{s}_t) = \mathcal{N}(\boldsymbol{\mu}_\theta(\boldsymbol{s}_t), \Sigma_\theta)$ is a local Gaussian policy with a local state-dependent mean and a global covariance matrix. $\boldsymbol{\mu}_\theta(\boldsymbol{s}_t)$ is a neural network with 3 hidden layers of sizes 100-50-25 and tanh nonlinearities at the first 2 layers, and Σ_θ is diagonal. For DDPG, the policy is deterministic and has the same architecture as $\boldsymbol{\mu}_\theta$ except that it has an additional tanh layer at the output. $V_\phi(\boldsymbol{s}_t)$ for baselines and GAE is fit with the same technique by Schulman et al. (2016), a variant of linear regression on Monte Carlo returns with soft-update constraint. For Q-Prop and DDPG, $Q_w(\boldsymbol{s}, \boldsymbol{a})$ is parametrized with a neural network with 2 hidden layers of size 100 and ReLU nonlinearity, where $\boldsymbol{a}$ is included after the first hidden layer.

Training details. This section describes parameters of the training algorithms and their hyperparameter search values in $\{\}$. The optimal performing hyperparameter results are reported. Policy gradient methods (VPG, TRPO, Q-Prop) used batch sizes of $\{1000, 5000, 25000\}$ time steps, step sizes of $\{0.1, 0.01, 0.001\}$ for the trust-region method, and base learning rates of $\{0.001, 0.0001\}$ with Adam (Kingma & Ba, 2014) for vanilla policy gradient methods. For Q-Prop and DDPG, Q_w is learned with the same technique as in DDPG (Lillicrap et al., 2016), using soft target networks with $\tau = 0.999$, a replay buffer of size 10^6 steps, a mini-batch size of 64, and a base learning rate of $\{0.001, 0.0001\}$ with Adam (Kingma & Ba, 2014). For Q-Prop we also tuned the relative ratio of gradient steps on the critic Q_w against the number of steps on the policy, in the range $\{0.1, 0.5, 1.0\}$, where 0.1 corresponds to 100 critic updates for every policy update if the batch size is 1000. For DDPG, we swept the reward scaling using $\{0.01, 0.1, 1.0\}$ as it is sensitive to this parameter.